arXiv:1611.09232v1 [stat.ML] 28 Nov 2016

Efficient Convolutional Auto-Encoding via Random Convexification and Frequency-Domain Minimization

Meshia C. Oveneke[1], Mitchel Aliosha-Perez[1], Yong Zhao[12], Dongmei Jiang[2], Hichem Sahli[13]

[1]VUB-NPU Joint AVSP Research Lab
Vrije Universiteit Brussel (VUB)
Deptartment of Electronics & Informatics (ETRO)
Pleinlaan 2, Brussels, Belgium
{mcovenek,maperezg,yzhao,hsahli}@etrovub.be
[2]VUB-NPU Joint AVSP Research Lab
Northwestern Polytechnical University (NPU)
Shaanxi Key Lab on Speech and Image Information Processing
Youyo Xilu 127, Xi'an 710072, China
jiangdm@nwpu.edu.cn
[3]Interuniversity Microelectronics Centre (IMEC)
Kapeldreef 75, 3001 Heverlee, Belgium

Abstract

The omnipresence of deep learning architectures such as deep *convolutional neural networks* (CNN)s is fueled by the synergistic combination of ever-increasing labeled datasets and specialized hardware. Despite the indisputable success, the reliance on huge amounts of labeled data and specialized hardware can be a limiting factor when approaching new applications. To help alleviating these limitations, we propose an efficient learning strategy for layer-wise unsupervised training of deep CNNs on conventional hardware in acceptable time. Our proposed strategy consists of randomly convexifying the *reconstruction contractive auto-encoding* (RCAE) learning objective and solving the resulting large-scale convex minimization problem in the frequency domain via *coordinate descent* (CD). The main advantages of our proposed learning strategy are: (1) single tunable optimization parameter; (2) fast and guaranteed convergence; (3) possibilities for full parallelization. Numerical experiments show that our proposed learning strategy scales (in the worst case) linearly with image size, number of filters and filter size.

1 Introduction

At the heart of the recent success of deep *convolutional neural networks* (CNN)s in several application domains such as computer vision, speech recognition and natural language processing, is the synergistic combination of ever-increasing *labeled datasets* and *specialized hardware* in the form of graphical processing units (GPU)s [16], field programmable gate arrays (FPGA)s [10] or computer clusters [18]. Despite this indisputable success, relying on huge amounts of labeled data and specialized hardware can be a limiting factor when approaching new applications. Furthermore, for applying deep learning techniques locally on platforms with limited resources, one cannot rely on huge amounts of labeled data and specialized hardware. We therefore argue that there is a growing need for resource-efficient *unsupervised learning strategies* capable of training deep CNNs on conventional hardware in acceptable time. The main purpose of unsupervised learning is to leverage the wealth of unlabeled data for disentangling the causal (generative) factors [2, 7]. In the context of supervised pattern classification, empirical evidence shows that: unsupervised pre-training (generative) helps in disentangling the class

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

manifolds in the lower layers of a deep CNNs, the upper layers are better disentangled when they are subject to supervised training (discriminative) and a combination of the two boosts the overall performance when the ratio of unlabeled to labeled samples is high [3, 4, 14].

In this work, we propose to train deep CNNs in a greedy layer-wise manner, using the *reconstruction contractive auto-encoding* (RCAE) learning objective [1]. The RCAE learning objective has been proven (theoretically and empirically) to capture the local shape of the data-generating distribution [2, 1], hence capturing the manifold structure of the data. Meanwhile, minimizing the RCAE learning objective involves solving a non-convex minimization problem, often addressed using *stochastic gradient descent* (SGD) methods [13]. Despite their empirical success when applied to deep CNNs, their key disadvantage is the computationally expensive manual tuning of optimization parameters such as learning rates and convergence criteria. Moreover, their inherently sequential nature makes them very difficult to parallelize using GPUs or distribute them using computer clusters [13]. To overcome the above-mentioned difficulties, we propose to convexify the RCAE learning objective. To this end, inspired by recent work such as [15, 8, 11], we propose to adopt a *random convexification* strategy by fixing the (non-linear) encoding parameters and only learning the (linear) decoding parameters. By further transforming the randomly convexified RCAE objective into the frequency domain using the *discrete Fourier transform* (DFT), we obtain a learning objective which we propose to minimize using *coordinate descent* (CD) [12, 17]. The main advantages of our proposed learning strategy are: (1) single tunable optimization parameter; (2) fast and guaranteed convergence; (3) possibilities for full parallelization.

2 Efficient Convolutional Auto-Encoding

The main motivation of this work is efficient unsupervised training of deep CNNs. To this end, we adopt a greedy layer-wise learning strategy and consider deep CNNs as a stack of single-layer CNNs with K filters and input space $\mathcal{X} \subset \mathbb{R}^{d \times d \times C}$, i.e. the space of C-channel $d \times d$ images [1] $\mathbf{x} = \left(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(C)}\right)$. We further propose to train each single-layer CNN using the *reconstruction contractive auto-encoding* (RCAE) learning objective [1]. The RCAE learning objective has been proven to capture the high-density regions of the (layer-wise) data-generating distribution for any twice-differential reconstruction (encoding-decoding) function and regularization parameter approaching 0 asymptotically. To this end, we consider the following parametric reconstruction function:

$$\begin{aligned} \mathbf{r}(\mathbf{x}; \boldsymbol{\theta}) &\triangleq \underbrace{\sum_{k=1}^{K} \mathbf{w}^{(k)}}_{\text{decoding}} * \underbrace{\mathbf{g}\Big(\sum_{c=1}^{C} \mathbf{a}^{(k)} * \mathbf{x}^{(c)} + \mathbf{b}^{(k)}\Big)}_{\text{encoding}} \\ &= \sum_{k=1}^{K} \mathbf{w}^{(k)} * \mathbf{h}^{(k)} \end{aligned} \tag{1}$$

with $\mathbf{g}(\cdot)$ denoting the entry-wise application of a twice-differentiable activation function $g : \mathbb{R} \to \mathbb{R}$. The model parameters $\boldsymbol{\theta}$ consist of $w \times w$ encoding filters $\{\mathbf{a}^{(k)}\}_{k=1}^{K}$, $h \times h$ encoding biases $\{\mathbf{b}^{(k)}\}_{k=1}^{K}$ and $w \times w$ decoding filters $\{\mathbf{w}^{(k)}\}_{k=1}^{K}$. We consider a valid convolution for the encoding function, yielding a dimension $h = d - w + 1$, and a full convolution for the decoding function.

The core idea behind our proposed layer-wise learning strategy is to sample each entry of the (non-linear) encoding parameters $\{\mathbf{a}^{(k)}\}_{k=1}^{K}$ and $\{\mathbf{b}^{(k)}\}_{k=1}^{K}$ i.i.d. from pre-determined density functions $p(a)$ and $p(b)$ respectively, and keeping them fixed while learning the (linear) decoding parameters $\{\mathbf{w}^{(k)}\}_{k=1}^{K}$ in the frequency domain. To this end, we define the complex-valued decoding parameters associated to $\{\mathbf{w}^{(k)}\}_{k=1}^{K}$ as $\{\mathbf{W}^{(k)}\}_{k=1}^{K}$, with $\mathbf{W}^{(k)} = \mathcal{F}\{\mathbf{w}^{(k)}\} \in \mathbb{C}^{d \times d}$ being the *discrete Fourier transform* (DFT). Given a set of N training images $\mathcal{D}_N = \{\mathbf{x}_n\}_{n=1}^{N}$ sampled i.i.d. from the data-

[1] Although we consider $d \times d$ images, the remaining calculations also hold for rectangular images.

generating distribution $p(\mathbf{x})$, we define the following frequency-domain RCAE objective:

$$\begin{aligned}
\mathcal{L}_{\text{RCAE}} &= \int_{\mathcal{X}} p(\mathbf{x}) \left[\|\mathbf{r}(\mathbf{x};\boldsymbol{\theta}) - \mathbf{x}\|_F^2 + \lambda \left\| \frac{\partial \mathbf{r}(\mathbf{x};\boldsymbol{\theta})}{\partial \mathbf{x}} \right\|_F^2 \right] d\mathbf{x} \\
&\approx \frac{1}{N} \sum_{n=1}^{N} \left[\|\mathbf{r}(\mathbf{x}_n;\boldsymbol{\theta}) - \mathbf{x}_n\|_F^2 + \lambda \left\| \frac{\partial \mathbf{r}(\mathbf{x};\boldsymbol{\theta})}{\partial \mathbf{x}} \Big|_{\mathbf{x}=\mathbf{x}_n} \right\|_F^2 \right] \\
&\propto \frac{1}{N} \sum_{n=1}^{N} \left[\|\mathcal{F}\{\mathbf{r}(\mathbf{x}_n;\boldsymbol{\theta}) - \mathbf{x}_n\}\|_F^2 + \lambda \left\| \mathcal{F}\left\{ \frac{\partial \mathbf{r}(\mathbf{x};\boldsymbol{\theta})}{\partial \mathbf{x}} \Big|_{\mathbf{x}=\mathbf{x}_n} \right\} \right\|_F^2 \right] \\
&= \frac{1}{N} \sum_{n=1}^{N} \left[\left\| \sum_{k=1}^{K} \mathbf{W}^{(k)} \odot \mathbf{H}_n^{(k)} - \mathbf{X}_n \right\|_F^2 + \lambda \left\| \sum_{k=1}^{K} \mathbf{W}^{(k)} \odot \mathbf{D}_n^{(k)} \right\|_F^2 \right]
\end{aligned} \tag{2}$$

where $\| \bullet \|_F$ is the Frobenius norm, $\odot$ denotes the Hadamard (entry-wise) product [5], $\mathbf{H}_n^{(k)} = \mathcal{F}\{\mathbf{h}_n^{(k)}\}$, $\mathbf{X}_n = \sum_{c=1}^{C} \mathcal{F}\{\mathbf{x}_n^{(c)}\}$, $\mathbf{D}_n^{(k)} = \mathbf{G}_n^{(k)} \odot \mathbf{A}^{(k)}$ with $\mathbf{G}_n^{(k)} = \mathcal{F}\left\{ \frac{\partial g(\mathbf{v})}{\partial \mathbf{v}} \Big|_{\mathbf{v}=\mathbf{v}_n^{(k)}} \right\}$ and $\mathbf{v}_n^{(k)} = \sum_{c=1}^{C} \mathbf{a}^{(k)} * \mathbf{x}_n^{(c)} + \mathbf{b}^{(k)}$. We used Parseval's theorem in conjunction with the convolution theorem [9] for transforming the randomly convexified RCAE objective from spatial to frequency domain. As a direct consequence, minimizing the transformed RCAE objective (2) reduces to solving d^2 independent K-dimensional regularized linear least-squares problems, which offers full parallelization possibilities. We propose to solve each of the independent K-dimensional problems using *coordinate descent* (CD) [17], which has recently witnessed a resurgence of interest in large-scale optimization problems due to its simplicity and fast convergence [12]. In the context of RCAE, the CD method consists of iteratively minimizing (2) along the k-th coordinate (filter), while keeping the remaining $K-1$ coordinates (filters) fixed, yielding the following strategy for updating the filters in the frequency domain:

$$\hat{\mathbf{W}}^{(k)} = \frac{\bar{\mathbf{H}}_N^{(k)} \odot \mathbf{X}_N - \sum_{i \neq k} \hat{\mathbf{W}}^{(i)} \odot \left(\mathbf{H}_N^{(i)} \odot \bar{\mathbf{H}}_N^{(k)} + \lambda \mathbf{D}_N^{(i)} \odot \bar{\mathbf{D}}_N^{(k)} \right)}{\mathbf{H}_N^{(k)} \odot \bar{\mathbf{H}}_N^{(k)} + \lambda \mathbf{D}_N^{(k)} \odot \bar{\mathbf{D}}_N^{(k)}} \tag{3}$$

where the bold fraction sign ($\dot{-}$) denotes the Hadamard (entry-wise) division and $\mathbf{H}_N^{(k)} = \sum_{n=1}^{N} \mathbf{H}_n^{(k)}$, $\mathbf{X}_N = \sum_{n=1}^{N} \mathbf{X}_n$ and $\mathbf{D}_N^{(k)} = \sum_{n=1}^{N} \mathbf{D}_n^{(k)}$. The main advantages of the derived filter update strategy (3) are: single tunable optimization parameter λ, fast and guaranteed convergence of $\hat{\mathbf{W}}^{(k)}$; possibilities for parallelization offered by the frequency-domain transformation of $\mathcal{L}_{\text{RCAE}}$. It is important to note that the incremental data acquisition involved in $\mathbf{H}_N^{(k)}$, $\mathbf{X}_N$ and $\mathbf{D}_N^{(k)}$ can also be computed in parallel. Once the decoding filters are learned in the frequency domain, they are transformed back to the spatial domain using the inverse DFT:

$$\hat{\mathbf{w}}^{(k)} = \mathcal{F}^{-1}\{\hat{\mathbf{W}}^{(k)}\} \tag{4}$$

At inference stage, we use the transpose of the learned decoding filters $\hat{\mathbf{w}}^{(k)}$ for computing k feature maps as follows: $\hat{\mathbf{h}}^{(k)} \triangleq \mathbf{g}(\sum_{c=1}^{C} \hat{\mathbf{w}}^{(k)T} * \mathbf{x}^{(c)})$.

3 Numerical Experiments

As a sanity check for the overall computational efficiency of our proposed learning strategy, we've implemented the CD-based frequency-domain RCAE minimization in MATLAB R2014a. We've used the built-in *fast Fourier transform* (FFT) for transforming the RCAE objective into the frequency domain. The hyperbolic tangent was used as activation function. The encoding filters and biases were randomly fixed by sampling their entries independently from a zero-mean normal distribution with standard deviations 0.1 and 0.01 respectively. For learning the decoding filters, we've implemented a fairly naive version of frequency-domain CD (3) by initializing each complex-valued filter as $\hat{\mathbf{W}}^{(k)} = \mathbf{0} + \mathrm{i}\mathbf{0}$ and only performing a single CD cycle trough the filter coordinates $k \in [1, K]$. For all the experiments, we used of the Caltech-256 Object Category dataset [6] and whitened the images.

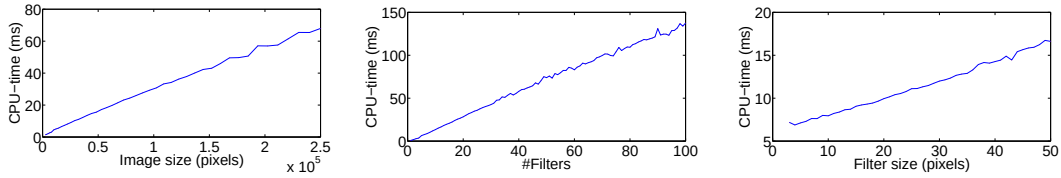

Figure 1: Time-complexity of CD-based RCAE minimization on Caltech-256 dataset. *Left*: CPU-time (in milliseconds) in function of image size (in pixels). *Middle*: CPU-time (in milliseconds) in function of number of filters. *Right*: CPU-time (in milliseconds) in function of filter size (in pixels).

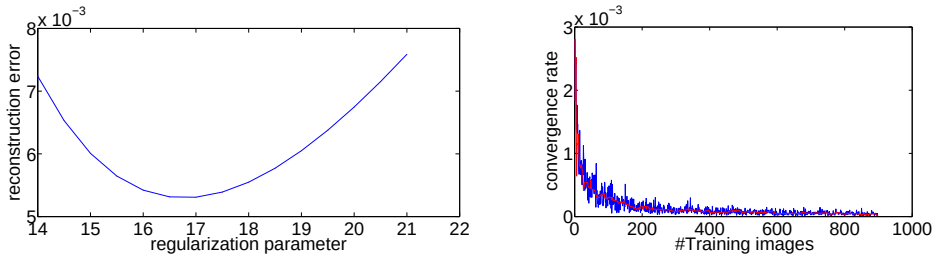

Figure 2: Reconstruction error and convergence rate of CD-based RCAE minimization on Caltech-256 dataset. *Left*: Reconstruction error in function of regularization parameter λ. *Right*: Convergence rate of the learned decoding filters in function of number of training images N. The convergence rate is computed using the average squared difference between consecutive filter updates. The original curve (blue) is noisy due to the average involved in the computation of the convergence rate. A smoothed curve (red) is therefore plotted on top, in order to highlight the general trend.

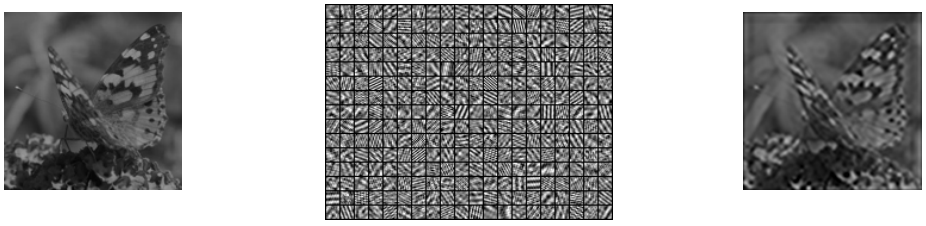

Figure 3: Visualization of learned decoding filters and reconstruction result on Caltech-256 dataset (after training with 400 images). *Left*: Previously unseen 244×244 testing image. *Middle*: 300 learned decoding filters. *Right*: Reconstruction result on previously unseen 244×244 testing image (left) using 300 learned decoding filters (middle).

As first experiment, we've measured the computational time-complexity in terms of CPU-time on an Intel® Core™ i7-2600 CPU @ 3.40 GHz machine. Figure (1) shows that our proposed method has a (worst-case) linear time-complexity w.r.t. image size, number of filters and filter size. Knowing that our naive implementation does not involve any particular form of parallelism, linear time is the best possible complexity we can achieve in situations where the algorithm has to read its entire input sequentially. In a second experiment, we study the influence of the regularization parameter λ on the reconstruction error and analyze the convergence rate of the learned decoding filters in function of the number training samples N. For robust estimation, the reconstruction error was averaged over a batch of 500 images that were not used during training. Figure (2), left, illustrates that the reconstruction error reaches a (global) minimum when the regularization parameter approaches $\lambda = 16.5$. Figure (2), right, illustrates that after roughly 400 training images, the learned decoding filters $\{\hat{\mathbf{w}}^{(k)}\}_{k=1}^{K}$ already settle. This clearly highlights the advantages of random convexification and frequency-domain minimization using CD. Figure (3) depicts the decoding filters and reconstruction result on a previously unseen testing image, obtained after CD-based minimization of the RCAE objective (2) using only 400 training images.

4 Conclusions

We've proposed an efficient learning strategy for layer-wise unsupervised training of deep CNNs. The main contributions of our proposed learning strategy are random convexification and frequency-domain transformation of the *reconstruction contractive auto-encoding* (RCAE) objective, which yields a relatively easy-to-solve large-scale regularized linear least-squares problem. We've proposed to solve this problem using *coordinate descent* (CD), with as main advantages: (1) single tunable optimization parameter; (2) fast and guaranteed convergence; (3) possibilities for full parallelization. Numerical experiments show that, with a fairly naive implementation, our proposed learning strategy scales (in the worst case) linearly with image size, number of filters and filter size. We also observe that, using relatively few training images, the learned filters already settle and yield decent reconstruction results on unseen testing images. We believe that the inherently parallel nature of our proposed learning strategy offers very interesting possibilities for further increasing the computational efficiency using more sophisticated implementation strategies.

Acknowledgments

This work is supported by the Agency for Innovation by Science and Technology in Flanders (IWT) – PhD grant nr. 131814, the VUB Interdisciplinary Research Program through the EMO-App project and the National Natural Science Foundation of China (grant 61273265).

References

[1] Guillaume Alain and Yoshua Bengio. What regularized auto-encoders learn from the data-generating distribution. *The Journal of Machine Learning Research*, 15(1):3563–3593, 2014.

[2] Yoshua Bengio, Aaron Courville, and Pierre Vincent. Representation learning: A review and new perspectives. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(8):1798–1828, 2013.

[3] P. P. Brahma, D. Wu, and Y. She. Why deep learning works: A manifold disentanglement perspective. *IEEE Transactions on Neural Networks and Learning Systems*, 27(10):1997–2008, Oct 2016.

[4] Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research*, 11(Feb):625–660, 2010.

[5] Gene H Golub and Charles F Van Loan. *Matrix computations*, volume 3. JHU Press, 2012.

[6] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. 2007.

[7] Irina Higgins, Loic Matthey, Xavier Glorot, Arka Pal, Benigno Uria, Charles Blundell, Shakir Mohamed, and Alexander Lerchner. Early visual concept learning with unsupervised deep learning. *arXiv preprint arXiv:1606.05579*, 2016.

[8] Guang-Bin Huang. What are extreme learning machines? filling the gap between frank rosenblatts dream and john von neumanns puzzle. *Cognitive Computation*, 7(3):263–278, 2015.

[9] David W Kammler. *A first course in Fourier analysis*. Cambridge University Press, 2007.

[10] Griffin Lacey, Graham W Taylor, and Shawki Areibi. Deep learning on fpgas: Past, present, and future. *arXiv preprint arXiv:1602.04283*, 2016.

[11] Xia Liu, Shaobo Lin, Jian Fang, and Zongben Xu. Is extreme learning machine feasible? a theoretical assessment (part i). *Neural Networks and Learning Systems, IEEE Transactions on*, 26(1):7–20, 2015.

[12] Yu Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.

[13] Jiquan Ngiam, Adam Coates, Ahbik Lahiri, Bobby Prochnow, Quoc V Le, and Andrew Y Ng. On optimization methods for deep learning. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 265–272, 2011.

[14] Tom Le Paine, Pooya Khorrami, Wei Han, and Thomas S Huang. An analysis of unsupervised pre-training in light of recent advances. *arXiv preprint arXiv:1412.6597*, 2014.

[15] Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *Advances in neural information processing systems*, pages 1313–1320, 2009.

[16] Rajat Raina, Anand Madhavan, and Andrew Y Ng. Large-scale deep unsupervised learning using graphics processors. In *Proceedings of the 26th annual international conference on machine learning*, pages 873–880. ACM, 2009.

[17] Stephen J Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, 2015.

[18] Ren Wu, Shengen Yan, Yi Shan, Qingqing Dang, and Gang Sun. Deep image: Scaling up image recognition.